

ABSTRACT

To avoid packet out-of-sequence problems, while providing good load balancing, each input port of a switch monitors the outstanding number of packets for each flow group. If there is an outstanding packet in the switch fabric, the following packets of the same flow group should follow the same path. If there is no outstanding packet of the same flow group in the switch fabric, the (first, and therefore subsequent) packets of the flow can choose a less congested path to improve load balancing performance without causing an out-of-sequence problem. To avoid HOL blocking without requiring too many queues, an input module may include two stages of buffers. The first buffer stage may be a virtual output queue (VOQ) and second buffer stage may be a virtual path queue (VPQ). At the first stage, the packets may be stored at the VOQs, and the HOL packet of each VOQ may be sent to the VPQ. By allowing each VOQ to send at most one packet to VPQ, HOL blocking can be mitigated dramatically.